

What is claimed is:

[Claim 1] 1. A system for use in providing data storage and data copies over a computer network, comprising:

a storage server system comprising one or more data storage servers that each comprise a data storage device and a network interface, each of said data storage servers operable to communicate over said network interface with at least one application client that will require data storage and at least one other data storage server; and

a data management system comprising at least one data management server operable to (a) define at least a first and a second cluster each comprising one or more data storage servers, (b) define at least one primary volume of data storage distributed over at least two of said storage servers within one of said clusters, said primary volume storing data from the application client, (c) define at least one remote volume of data storage distributed over one or more of said storage servers within one of said clusters; (d) create snapshots of said primary volume; and (e) copy data from said snapshots over the computer network to said remote volume.

[Claim 2] 2. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein each of said snapshots provides a view of the data stored at said primary volume at the point in time of said snapshot.

[Claim 3] 3. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein an application client is operable to read data stored in said snapshots at said primary volume.

[Claim 4] 4. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein an application client is operable to read data stored in said snapshots at said remote volume.

[Claim 5] 5. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein each snapshot includes data that has been modified at said primary volume since a previous snapshot of said primary volume.

[Claim 6] 6. The system for use in providing data storage and data copies over a computer network, as claimed in claim 5, wherein said snapshots are copied to remote snapshots associated with said remote volume.

[Claim 7] 7. The system for use in providing data storage and data copies over a computer network, as claimed in claim 5, wherein said snapshots are copied from said primary volume to said remote volume and at least a second remote volume distributed over one or more of said storage servers within one of said clusters.

[Claim 8] 8. The system for use in providing data storage and data copies over a computer network, as claimed in claim 5, wherein said snapshots are copied from said primary volume to said remote volume and at least a second remote volume distributed over one or more of said storage servers within one of said clusters, and wherein the source of said snapshots copied to said second remote volume is selected based on at least one of the volume most likely to be available, the least loaded volume, the volume with the highest bandwidth connection to the network, and the volume with a least costly connection to the network.

[Claim 9] 9. The system for use in providing data storage and data copies over a computer network, as claimed in claim 5, wherein said snapshots are copied from said primary volume to said remote volume and are copied from said remote volume to a second remote volume distributed over one or more of said storage servers within one of said clusters.

[Claim 10] 10. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said snapshots are created according to a predetermined schedule defined by said data management system.

[Claim 11] 11. The system for use in providing data storage and data copies over a computer network, as claimed in claim 10, wherein said snapshots are copied to remote snapshots associated with said remote volume according to said predetermined schedule.

[Claim 12] 12. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said data management system is further operable to designate said primary volume as a second remote volume that is not able to write data from application clients.

[Claim 13] 13. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said data management system is further operable to designate said remote volume as a second primary volume, said second primary volume storing data from at least one application client independently of said primary volume.

[Claim 14] 14. The system for use in providing data storage and data copies over a computer network, as claimed in claim 13, wherein said remote volume is designated as said second primary volume following a failure of said primary volume.

[Claim 15] 15. The system for use in providing data storage and data copies over a computer network, as claimed in claim 13, wherein said remote volume is designated as said second primary volume following a determination by a user to create a second primary volume.

[Claim 16] 16. The system for use in providing data storage and data copies over a computer network, as claimed in claim 13, wherein said data management system is

further operable to designate said primary volume as a second remote volume that is not able to write data from application clients.

[Claim 17] 17. The system for use in providing data storage and data copies over a computer network, as claimed in claim 16, wherein said data management system is operable to copy data from a snapshot of said second primary volume to said second remote volume.

[Claim 18] 18. The system for use in providing data storage and data copies over a computer network, as claimed in claim 17, wherein said data management system is operable to generate a snapshot of said primary volume prior to designating said primary volume as said second remote volume.

[Claim 19] 19. The system for use in providing data storage and data copies over a computer network, as claimed in claim 18, wherein said data management system is operable to resynchronize said primary volume with said second primary volume.

[Claim 20] 20. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said primary volume comprises a plurality of logical blocks of data.

[Claim 21] 21. The system for use in providing data storage and data copies over a computer network, as claimed in claim 20, wherein each of said plurality of logical blocks of data comprises a plurality of physical blocks of data, each physical block of data comprising a unique physical address associated with said data storage device and data to be stored at said unique physical address.

[Claim 22] 22. The system for use in providing data storage and data copies over a computer network, as claimed in claim 20, wherein said snapshots comprise pointers to logical blocks of data stored at said cluster.

[Claim 23] 23. The system for use in providing data storage and data copies over a computer network, as claimed in claim 20, wherein

wherein each of said logical blocks of data are copied from said primary volume to said remote volume and at least a second remote volume distributed over one or more of said storage servers within one of said clusters, and wherein the source of each of said logical blocks of data copied to said second remote volume is selected based on at least one of the volume most likely to be available, the least loaded volume, the volume with the highest bandwidth connection to the network, and the volume with a least costly connection to the network.

[Claim 24] 24. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said network interface is adapted to connect to one of an Ethernet network, a fibre channel network, and an infiniband network.

[Claim 25] 25. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said data management system is

operable to copy data from said snapshots to said remote volume independently of network protocol.

[Claim 26] 26. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said data management system is operable to copy data from said snapshots to said remote volume independently of network link bandwidth.

[Claim 27] 27. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said data management system is operable to copy data from said snapshots to said remote volume independently of network latency.

[Claim 28] 28. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said data management system is operable to copy data from said snapshots to said remote volume at a selected maximum bandwidth.

[Claim 29] 29. The system for use in providing data storage and data copies over a computer network, as claimed in claim 28, wherein said selected maximum bandwidth is adaptively set based on the network bandwidth capacity and utilization of the network.

[Claim 30] 30. The system for use in providing data storage and data copies over a computer network, as claimed in claim 28, wherein said selected maximum bandwidth is adjusted based on time of day.

[Claim 31] 31. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said first primary volume is located at a first cluster and said first remote volume is located at a second cluster.

[Claim 32] 32. The system for use in providing data storage and data copies over a computer network, as claimed in claim 31, wherein said first cluster and said second cluster are located at different geographic locations.

[Claim 33] 33. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said data management server is a distributed data management server distributed over one or more data storage servers.

[Claim 34] 34. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said data management server is further operable to redefine said primary volume to be distributed over one or more data storage servers that are different than said at least two data storage servers while copying data from said snapshots over the computer network to said remote volume.

[Claim 35] 35. The system for use in providing data storage and data copies over a computer network, as claimed in claim 1, wherein said data management server is

further operable to define at least one replica volume of data storage distributed over one or more of said data storage servers within one of said clusters, said replica volume storing data stored at said primary volume.

[Claim 36] 36. The system for use in providing data storage and data copies over a computer network, as claimed in claim 35, wherein said data management server is operable to create snapshots of said replica volume corresponding to said snapshots of said primary volume, and wherein the source of said snapshots copied to said remote volume selected based on at least one of the volume most likely to be available, the least loaded volume, the volume with the highest bandwidth connection to the network, and the volume with a least costly connection to the network.

[Claim 37] 37. The system for use in providing data storage and data copies over a computer network, as claimed in claim 35, wherein in the event of a failure associated with said primary volume, said data management server is operable to copy said snapshots from said replica volume to said remote volume.

[Claim 38] 38. The system for use in providing data storage and data copies over a computer network, as claimed in claim 37, wherein said failure is at least one of a data storage server failure and a network failure.

[Claim 39] 39. A method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, comprising:

defining a first primary volume of data storage distributed over at least two data storage servers within a first cluster of data storage servers;
generating a first primary snapshot of said first primary volume, said first primary snapshot providing a view of data stored at said first primary volume at the time said first primary snapshot is generated;
creating a first remote volume distributed over one or more data storage servers within a cluster of data storage servers;
linking said first remote volume to said first primary volume; and
copying data from said first primary snapshot to a first remote snapshot associated with said first remote volume.

[Claim 40] 40. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 39, further comprising:

generating a second primary snapshot of said first primary volume, said second primary snapshot providing a view of data stored at said first primary volume at the time said second primary snapshot is generated; and
copying data from said second primary snapshot to a second remote snapshot associated with said first remote volume.

[Claim 41] 41. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in

claim 39, wherein said second primary snapshot includes data that has been modified at said first primary volume since said step of generating a first primary snapshot.

[Claim 42] 42. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 39, further comprising:

copying data from said first snapshot to a second remote volume distributed over one or more storage servers within a cluster of data storage servers.

[Claim 43] 43. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 39, further comprising:

copying said first remote snapshot from said first remote volume to a second remote volume distributed over one or more storage servers within a cluster of data storage servers.

[Claim 44] 44. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 39, wherein said steps of generating first and second primary snapshots are performed according to a predetermined schedule defined by a data management system.

[Claim 45] 45. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 39, wherein said steps of copying said first and second primary snapshots to said first and second remote snapshots are performed according to a predetermined schedule defined by a data management system.

[Claim 46] 46. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 39, further comprising:

designating said first remote volume as a second primary volume, said second primary volume storing data from at least one application client independently of said first primary volume.

[Claim 47] 47. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 46, wherein said step of designating is performed following a failure of said first primary volume.

[Claim 48] 48. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 47, wherein said step of designating is performed following a determination by a user to create a second primary volume.

[Claim 49] 49. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 48, further comprising:

designating said first primary volume as a second remote volume that is not able to write data from application clients.

[Claim 50] 50. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 49, further comprising:

copying data written to said second primary volume to said second remote volume.

[Claim 51] 51. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 49, wherein said step of designating said first primary volume as a second remote volume comprises:

generating a snapshot of said first primary volume; and
designating said first primary volume as said second remote volume.

[Claim 52] 52. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 49, further comprising resynchronizing said second primary volume with said second remote volume.

[Claim 53] 53. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 52, wherein said step of resynchronizing comprises:

generating a second primary snapshot of said second primary volume providing a view of data stored at said second primary volume at the time said second primary snapshot is generated;

generating a second remote snapshot of said second remote volume providing a view of data stored at said first primary volume at the time said third primary snapshot is generated;

copying data that has been modified at said second primary volume to said second remote volume.

[Claim 54] 54. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 39, wherein said step of creating a first remote volume comprises:

creating a volume at a cluster of data storage servers;

designating said volume as a remote volume;

linking said remote volume to said first primary volume; and

setting a maximum bandwidth at which data may be copied to said remote volume.

[Claim 55] 55. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 54, wherein said step of setting is based on network bandwidth capacity and network utilization.

[Claim 56] 56. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 54, wherein said step of setting comprises:

scheduling a maximum bandwidth at which data may be copied to said remote volume.

[Claim 57] 57. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 56, wherein said step of scheduling is based on at least one of time of day and day of the week.

[Claim 58] 58. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 39, wherein said data management system is a distributed data management server distributed over one or more of said data storage servers.

[Claim 59] 59. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 39, wherein said primary volume comprises a plurality of logical blocks of data, and wherein said step of generating a first primary snapshot comprises moving a pointer associated with each of said plurality of logical blocks of data from said primary volume to said first primary snapshot.

[Claim 60] 60. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 39, wherein said step of copying comprises:

copying a first portion of said first primary snapshot to said first remote snapshot;
recording that said first portion has been copied; and
copying a second portion of said first primary snapshot to said first remote snapshot.

[Claim 61] 61. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 60, wherein said step of copying a second portion is interrupted, and said step of copying a second portion is re-started based on said step of recording.

[Claim 62] 62. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in claim 60, wherein the amount of data included in said first portion is based on an amount of data contained in said first primary snapshot.

[Claim 63] 63. The method for copying data from a primary data storage volume to a remote data storage volume in a distributed data storage system, as claimed in

claim 60, wherein the amount of data included in said first portion is determined based on an elapsed time period since starting said step of copying a first portion.